DOI: 10.1111/phpe.12182

ORIGINAL ARTICLE

PHILOSOPHICAL PERSPECTIVES WILEY

Compositionality and constituent structure in the analogue mind

Sam Clarke

University of Southern California

Correspondence
Sam Clarke, University of Southern California.
Email: Sam.clarke@usc.edu

Abstract
I argue that analogue mental representations possess a canonical decomposition into privileged constituents from which they compose. I motivate this suggestion, and rebut arguments to the contrary, through reflection on the *approximate number system*, whose representations are widely expected to have an analogue format. I then argue that arguments for the compositionality and constituent structure of these analogue representations generalize to other analogue mental representations posited in the human mind, such as those in early vision and visual imagery.

1 | INTRODUCTION

A familiar view in philosophy and cognitive science is that (certain) mental representations comprise *compositionally structured constituents*. Like sentences of natural language, they are *compositionally structured* in that the meaning of a (complex) mental representation is a function of the meaning of its parts plus their mode of combination. Moreover, such parts qualify as privileged *constituents*, in the sense that they contribute to the meaning of the complex representations they comprise and do so in a way that other representation parts do not. So, just as the sentence 'John loves Mary' (understood to mean, roughly, that *John loves Mary*, and parsed [((John)$_{NP}$ ((loves)$_{V}$ (Mary)$_{NP}$) $_{VP}$)]) has a canonical decomposition into privileged constituents from which its meaning is built – i.e., 'John' and 'loves Mary', but not 'John loves' or 'John…Mary' – similar points are meant to apply to mental states, like my belief that *John loves Mary* (Fodor 2008). For brevity, let's call the view that complex mental states comprise compositionally structured constituents *Privileged Compositionality*.

How far does Privileged Compositionality extend? In the above example (my belief that *John loves Mary*), proponents of Privileged Compositionality will typically hold that the representational vehicles involved bear arbitrary similarities and differences to the entities or properties represented. Just as the word 'John' need not resemble John in any obvious or semantically relevant way, the neuropsychological vehicle that realizes my concept JOHN (and is tokened in the above belief) need not resemble John, in anything like the way a picture of John might (c.f. Greenberg 2013). In this way, proponents of Privileged Compositionality typically hold that the mental representations they describe are sentence-like in a further respect, comprising compositionally structured *digital* constituents. But what about the myriad *non-digital* mental representations that cognitive scientists find reason to posit? For instance, what about visual images (Shepard 1978; Kosslyn 1980; Johnson-Laird 2001), early visual states (Neisser 1967; Treisman 1986; Clarke 2022a), or the representations of 'core cognition' (Carey 2009; Beck 2015; Clarke 2022b) which are often said to be analogue, iconic, or otherwise picture-like? Are these representations composed from privileged constituents, or are they beyond the remit of Privileged Compositionality?

Many researchers seem to suggest the latter. For instance, Beck (2012) critiques the view that analogue magnitude representations can or could systematically recombine. This is striking, since systematic recombinability is something that the compositionality of these representations could seem to imply (Fodor & McLaughlin 1990). Meanwhile, Fodor (2007) follows Kosslyn (1980) in holding that while the iconic, or analogue, representations involved in visual imagery and early vision are compositional, they are characterized by their conformity to *the picture principle*. Accordingly, these representations are considered analogous to photographs in that *every* part of the vehicle represents some part of the scene depicted: just as you can (allegedly) "Take a picture of a person, [and] cut it into parts however you like" yet never be left with a part which fails to picture some part of the person or scene originally depicted, Fodor proposes that analogue or iconic mental representations "have interpretable parts, but they don't have constituents. Or, if you prefer, all the parts [of some such representation] is among its constituents" (2007: 108). Indeed, Fodor takes this to be *the* fundamental difference between analogue or iconic representations and the digital, discursive, or linguiform representations he posits in conceptual thought. This is striking, for while Fodor leaves "open that some kinds of representations are neither iconic nor discursive", he claims that he "cannot think of a good candidate" (107). He thus seems to suggest that *all* the non-digital mental representations we have grounds to posit lack privileged constituents entirely.

In the present treatment, I'll argue that these conclusions are mistaken; the human mind is, in fact, rife with compositionally structured representations which comprise privileged *analogue* constituents. To develop this point, I'll consider a scientifically important case study: approximate number representations. In Section 2, I'll introduce this case study in some detail. But first, I wish to note that my focus on approximate number representations is non-arbitrary. Researchers who study these representations converge on the view that they have a non-digital *analogue* format (e.g., Beck, 2015; Carey 2009; Condry & Spelke, 2008; Clarke 2022b; Dehaene, 2011; Feigenson, Carey & Hauser, 2002; Feigenson, Dehaene & Spelke, 2004; Lipton & Spelke, 2003; Nieder & Miller 2004; Slaughter, Kamppi & Paynter, 2006). So, unlike Arabic numerals, or constituents in a language of thought (e.g., my concept *JOHN*), researchers who study approximate number representations typically accept that these representations represent numbers via an isomorphism between content-bearing properties of their vehicle and the numbers represented (a bit like how an analogue thermometer represents the temperature by having its mercury-level *mirror* this). This suggestion will be spelled out in Section 4. What's important is that various proponents of this suggestion have offered seductive (and connected) arguments for the view that analogue number representations could not systematically recombine (which, as mentioned, could seem

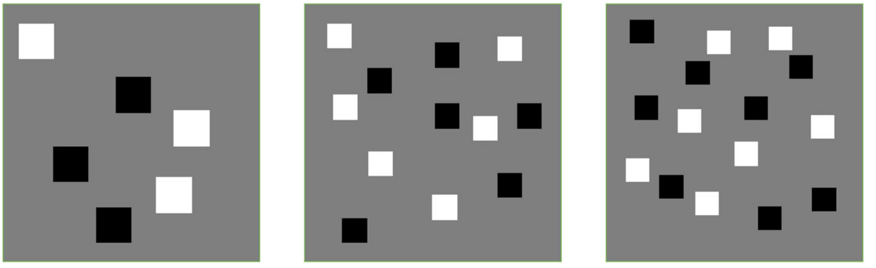

FIGURE 1 It is easier to see that the middle box contains more dots that the left-hand box, than it is to see that the right-hand box contains more dots than the middle box.

to undermine their compositionality), and for the view that their analogue format implies their conformity to the abovementioned picture principle (which is typically taken to preclude their privileged constituency). So, while it often seems to be tacitly assumed that analogue representations lack constituents and/or compositional structure, theorists studying approximate number representations are interesting in that they offer arguments to this effect. This makes them an ideal stalking horse in what follows.

2 | A PHILOSOPHER'S GUIDE TO APPROXIMATE NUMBER REPRESENTATION

Since it will be helpful to have a basic grasp of the approximate number representations I'll discuss, I want to begin by introducing these representations in some detail. This will lay the groundwork for me to then explain how and why they compose (Section 3); why they have an analogue format; why this still involves them comprising privileged constituents (Section 4); and why these points apply to other analogue representations that have been posited in the human mind (Section 5).

2.1 | The Approximate Number System

The approximate number system (ANS) is a well-studied psychological system (or collection of systems) which enables organisms to represent the approximate number of items in an observed collection, quickly and without recourse to explicit counting (Clarke & Beck 2021a). There is much that could be said about this system and the approximate number representations it produces. What matters for our purposes is that while the ANS facilitates numerical discriminations among relatively large numbers of items, these numerical discriminations are distinctive in their conformity to Weber's Law. Thus, the reliability with which the ANS discriminates two numbers is predicted by the ratio between these, rather than their absolute difference (the further from 1:1 the better).

To illustrate, consider the images below (Figure 1). Each image contains a different number of dots. Yet you will probably find it easier to quickly identify – without counting and simply by looking – that the middle image contains more dots than the left-hand image than you will identifying that the right-hand image contains more dots than the middle image. Crudely, this is because the ratio between the number of dots in the left-hand image and the central image is further from 1:1. Thus, your ANS finds it easier to facilitate the former discrimination. (Your ability to discriminate the number of dots is not due to their total surface area or brightness since these non-numerical properties are equated across all three images – but see Yousif & Keil [2021] for some interesting complications.)

2.2 | Evidence for an ANS

The preceding example is quick and dirty. However, the suggestion that humans quickly discriminate large numbers of items, approximately, and in accord with Weber's Law, is well-studied and well-established. In one study, Barth and colleagues (2003) presented adult humans with displays of dots similar to those in Figure 1. Despite short presentation times (too short for participants to explicitly count the number of dots in the displays) participants could identify a match in the number of items they contained, albeit imperfectly and with errors predicted by Weber's Law. Indeed, they did so in both intra- and inter-modal conditions, as when they were asked to compare the number of dots on a screen with the number of heard tones in an acoustic sequence. This is significant since inter-modal comparisons naturally eliminate non-numerical confounds in these comparisons. For instance: While the examples in Figure 1 are matched in total surface area and brightness, critics will sometimes argue that discriminations could still be based on other non-numerical properties, such as the spatial density or convex hull of the dots (Gebuis et al. 2016; Leibovich et al. 2017). The point to note is that heard tones do not possess a surface area, brightness, convex hull, *or* spatial density. Thus, it is not plausible that participants in Barth et al.'s study based their comparisons on these non-numerical properties. Instead, it seems they based their comparisons on the perceived *number* of items in each collection.

Perhaps this falls short of establishing that participants were relying on a distinct psychological system (or cluster of systems) for approximate number representation (an ANS). After all, the adults tested in this study had endured a formal math education. Thus, you might suspect that they still somehow managed to count the dots (albeit imprecisely) using the domain general cognitive resources involved in their flexible thinking about numbers, despite the brief presentation times employed. What's important to note is that there are many reasons to doubt this. For one, performance in these tasks is largely unaffected when participants perform a verbal shadowing task designed to consume verbal working memory and tax domain general resources (e.g., repeating a word, or poem, throughout the experiment) (Cordes et al. 2001). Meanwhile, performance *is* influenced by myriad judgement-independent illusions. For instance, circles that are connected with thin lines (effectively turning pairs of circles into single dumbbell-shaped objects) are irresistibly perceived to be less numerous than otherwise identical arrays in which lines are free floating (He et al. 2009) or in which the connecting lines contain small breaks (Franconerri et al. 2009). Indeed, this persists even when subjects try to ignore the lines completely and focus only on the circles (see also: Kirjakovski & Matsumoto 2016). Conversely, collections of items which cohere in color (Qu et al. 2022) and/or orientation (DeWind et al. 2020) are irresistibly perceived to be more numerous than collections which don't, *ceteris paribus*. Critically, these illusions persist even when participants know what's happening, with the effects influencing the Weber fraction they can reliably discriminate. Thus, displays continue to *look* more or less numerous, even when subjects know that they are undergoing an illusion. This shows that explicit judgements about number dissociate from the conclusions of the ANS, indicating that largely independent psychological processes and representations are involved (for my preferred account of the architectural relationship between these processes and representations, see Clarke 2021).

2.3 | Developmental Studies

One of the more exciting things about the ANS is that it is not simply found in adult humans – it is also found in a wide range of non-human animals (Dehaene 2011, ch.1) and young human infants (even newborns, under 5 days old – Izard et al. 2007). For instance, in a much-celebrated study, Xu and Spelke (2000) habituated six-month-old infants to visual arrays containing either 8 or 16 dots. When infants were habituated to an 8-dot array, they recovered interest when presented with a 16- or 4-dot array, but not a 12-dot array. Meanwhile, infants who were habituated to a 16-dot array dishabituated to a 32- or 8-dot array, but not a 12- or 24-dot array. Since confounding variables such as brightness, density, and dot size were controlled for, these findings were interpreted as showing that the six-month-old infants could represent and discriminate the approximate number of items in the two collections provided they differ by a suitably large ratio (e.g., 1:2). Subsequent studies then showed that these discriminative capacities improve with age. For instance, nine-month-olds were found to reliably discriminate collections that differ by a ratio of just 2:3 under comparable conditions (Lipton & Spelke 2003). In either case, the discriminations were ratio sensitive, thereby implicating psychological mechanisms with the distinctive profile of an ANS.

Of course, infants are a nightmare to study (often bursting into tears, or falling asleep, mid-experiment). Moreover, they can't talk, so their cognitive capacities can only be measured via implicit measures like looking time. This leads some researchers to question the reliability of these results (Leibovich et al. 2017). However, Libertus and Brannon (2010) replicated these findings using a novel change-detection paradigm and even showed that individual differences in task performance were reliable from one testing session to the next.

Indeed, similar results obtain in older children who can talk yet remain unable to count. In a particularly elegant study, Barth and colleagues (2005) presented 5year-old children with a collection of blue dots on a screen, before occluding these dots from view with an opaque block. At this point, a second collection of (red) dots appeared on the screen. The children were then asked a simple question: *Are there more blue dots* (behind the block) *or red dots* (on the screen)*?* The children could answer this question correctly, despite remaining unable to count properly. However, their accuracy was predicted by the ratio among the collections. Thus, children were more accurate in saying if there were more blue dots or red dots when there was a 4:7 ratio of blue to red dots than when there was a 4:6 ratio, and they were more accurate when there was a 4:6 ratio of blue to red dots than when there was a 4:5 ratio. In this way, their performance was (again) predicted by Weber's Law. In fact, these results persisted unaffected in an intermodal task with the children performing *as well* when the second collection of dots was replaced with a sequence of heard tones, and they were asked to report whether there were more blue dots in the original collection or more heard tones in the sequence. For reasons noted, this is significant since intermodal comparisons naturally eliminate non-numerical confounds, indicating that performance pertained to *the number* of items in the collections/sequences.

2.4 | Rich computations

Finally, it is worth stressing that while the abovementioned experiments simply involved participants making ordinal judgements, or same-different comparisons, the ANS facilitates more sophisticated numerical computations also. For instance, in a further experiment from Barth et al.'s (2005), five-year-old children saw a collection of blue dots appear and move behind an opaque block, at which point a second collection of blue dots appeared and did the same, such that the children could no longer see either collection. A collection of red dots then appeared, and the children were asked if there were more blue dots in total, behind the block, or more red dots on the screen. They were significantly above chance in answering this question correctly, indicating that they could *add* the approximate numbers of blue dots in the first two collections together, and then compare their total to that of the red dots (see also McCrink & Wynn 2004). A follow up study showed that children could also use their ANS to perform *subtraction* operations, as when dots were removed from the collection hidden behind the block and they were asked whether the dots left behind the block were more or less numerous than dots in a separate collection (Barth et al. 2006). Qu and colleagues (2021) recently went further, finding that children of a similar age group can even use their ANS to perform approximate multiplications (see also McCrink & Spelke 2010), with Szkudlarek and colleagues (2022) also finding a related capacity for division (see also McCrink & Spelke 2016).

What's notable about these latter studies is that the computations involved cannot simply be sensitive to the fact that one collection of dots contains *more* items than another. After all, a bare representation of *more/less* (or *same/different*) would not enable subjects to (e.g.,) add the quantity of items in two collections together and then compare their total to that of some separate collection (is *less number* + *less number* supposed to be bigger or smaller than *more number??*). Rather, they seem to require that participants' ANSs were representing relatively determinate or specific numerical quantities in the observed collections, such that these relatively determinate quantities could be added, subtracted, multiplied, or otherwise processed. Of course, this leaves open the vexed issue of how these (relatively determinate) numerical quantities are being represented. For everything that has been said here, participants could be representing precise quantities somewhat inaccurately (e.g., sometimes representing there to be precisely 5 dots when presented with 6) or representing these precise quantities imprecisely (e.g., representing there to be 6ish dots when presented with 6 – Lyons 2021). Alternatively, they could involve subjects representing some numerical range (e.g., 5–7 dots – Ball 2017) or (on my preferred view) a numerical range with some probability distribution or confidence level attached (Halberda 2016). All these possibilities can accommodate the abovementioned results provided that (e.g.,) noise in the extraction of number suffices to explain the ANS's characteristic imprecision and conformity to Weber's Law (Clarke & Beck 2021b). It is simply for ease of explication, that I will proceed by talking as though the ANS represents precise natural numbers (e.g., 6, 7, 8...) except when explaining why my arguments do not turn on this assumption.[1]

3 | COMPOSITIONALITY IN THE ANS

Having introduced the ANS, and the approximate number representations it produces, I'll now turn to my suggestion that approximate number representations exhibit Privileged Compositionality, featuring as, and comprising, privileged *analogue* constituents in compositionally structured representations. I'll begin by arguing that approximate number representations compose.

To remind readers, compositionality obtains when the meaning of a complex representation is a function of what its parts mean plus the mode by which they are combined. Put differently: parts of a compositionally structured representation make uniform semantic contributions to the complex representations they comprise, such that the meaning of these complex representations is determined by the (fixed) semantic contributions these parts make, plus the structural or syntactic relations that hold between them.

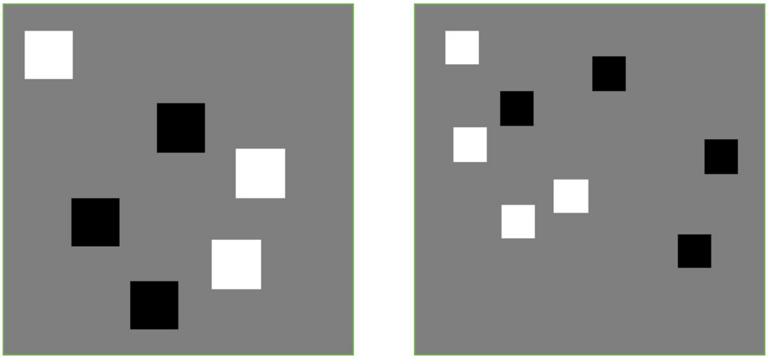

FIGURE 2 Assuming perfect veridicality, we might suppose that a glance at this Figure will result in my ANS representing that the left-hand box contains 6 (or 6ish) dots, that the right-hand box contains 8 (or 8ish) and that 8 (or 8ish) is bigger than 6 (or 6ish).

I offer three reasons to think approximate number representations compose in this way:

(i) They seem to
(ii) The most prominent challenge to this suggestion fails
(iii) The specific way in which this challenge fails is most straightforwardly explained by the compositionality of these representations

3.1 | Reason 1: Approximate number representations *seem* to compose

An initial reason to think that approximate number representations feature in compositionally structured representations is simply that they *seem* to feature in compositionally structured representations. Consider two examples.

First, approximate number representations seem to feature in compositionally structured representations when they underlie our ordinal (or same/different) judgements about the number of items in two collections. Take Figure 2. If one stops to count, they will find that there are 6 dots in the left-hand box and 8 dots in the right. Thus, (in the good case/assuming perfect accuracy) a glance at the figure might result in my ANS representing the left-hand collection of dots as 6 (or 6ish) in number and the right-hand collection as 8 (or 8ish) in number (after all, we have just seen that approximate number representations represent relatively determinate numerical quantities, and not just *more* or *less*, etc.). In addition, it seems that my ANS will thereby come to represent the left-hand collection as smaller in number than the right-hand collection. After all, children (who lack the capacity to count precisely) use their ANS to correctly identify the bigger of two analogous collections under analogous conditions (e.g., Barth et al. 2005). Thus, we might suppose that the total content of the representation produced is something like *the right-hand collection containing 8 is bigger in number than the left-hand collection containing 6*.

The point to note is that *prima facie*, the semantic value of this complex state decomposes into (at least) two distinct approximate number representations (meaning *that collection is 8 in number* and *that collection is 6 in number* respectively). In either case, these representations concern independent collections and seem to represent them independently of the other. For instance, if the right-hand box came to contain 10 dots because two additional dots were added to the collection, my ANS might come to represent that *the right-hand collection containing 10 is bigger in number than the left-hand collection containing 6*. But, in this scenario, we would expect the part of the representation responsible for representing the left-hand collection as 6 in number, to continue making the same semantic contribution to the complex state in which it features.

Second, similar points apply to individual approximate number representations when they represent some specific collection as 6 (or 6ish, etc.) in number. When they do, these representations seem to comprise a demonstrative element (compare Pylyshyn 2007) or representation that picks out the collection (compare Echeverri 2017) *and* a syntactically independent numerical element which attributes some numerical value to this. This allows these distinct elements to then compose in the manner of an object file, such that the total content of the representation produced is something like *[that collection]$_{subject}$ [is 6 (or 6ish) in number]$_{predicate}$* (see: Feigenson 2011; Green & Quilty-Dunn 2020).[2]

One reason to say this is that the identification of a collection, apt for enumeration, flexibly supports various non-numeric computations also. For instance, Odic and Halberda (2015) found that when subjects were asked to identify which of two collections had a larger total surface area this would affect their looking behavior (prompting fewer and longer saccades between targets) as compared with a condition in which they were asked to identify the more numerous of the collections. This suggests that the collection was picked out prior to and independently of the (numerical and/or non-numerical) magnitudes attributed to it, such that it could be scanned and processed in a task relevant manner (see also Odic 2018). Conversely, we have considered cases in which a collection of dots a^1 is added to an existing collection *A*, and the resulting number of items in collection *A* is then compared to that of a separate collection *B* (Barth et al. 2005). Here, the numerical quantity of collection a^1 seems to be added to collection *A*, such that our ANS updates its description of *A* (enabling it to then be compared to *B*). However, it is unlikely that this involves us retaining reference to collection a^1 as a subset of *A*, not least because the number of numerical ensembles that can be retained in working memory is extremely limited (Halberda et al. 2006; Zosh et al. 2007). Thus, when the ANS represents a collection as 6 in number, it seems likely that the *collection identifying* element of the representation functions independently of its *numerical attribute* and vice versa. Yet when these elements are combined – and we thereby come to represent, e.g., that ([*that collection*]$_{\text{subject}}$ [*is 6 in number*]$_{\text{predicate}}$) – the meaning of this complex representation would seem to be a function of what these elements mean plus the mode by which they are combined (i.e., with the numerical element predicated of a specific collection).

3.2 | Reason 2: Arguments to the contrary fail

Despite enjoying *prima facie* plausibility, the claim that approximate number representations feature in compositionally structured representations of the above sort has not gone unchallenged. To see why, notice that (given the compositionality of the complex representations described above) approximate number representations should *systematically recombine*. Thus, the competence to represent that *8>6*, and the ability to represent that *7>5*, should imply the competence to represent that *8>7*. For if approximate number representations (of e.g., *5, 6, 7, 8…*) make fixed semantic contributions to the complex states in which they feature, such that the meaning of these complex states is a function of what these (and any other relevant) representation parts mean plus the mode by which they are combined, then the competence to deploy an approximate number representation of (e.g.,) 7 in one context should imply the competence to deploy it in another, at least when the structural relations among the relevant representation parts remain the same. Indeed, systematic competences of this sort could be seen to follow from these representations' compositionality as a matter of "nomological necessity" (Fodor & McLaughlin 1990: 188).

The trouble is: Approximate number representations have been argued not to systematically recombine. In particular, Beck (2012) rejects their systematic recombinability on the grounds that the discriminability of two approximate number representations is a function of the ratio between the numerical values represented (the further from 1:1, the better).[3]

To illustrate, consider the 6month-old infants tested by Xu and Spelke (2000). As discussed in Section 2, these infants were habituated to collections, containing either 8 or 16 items. When habituated to 8 items, the infants were found to reliably dishabituate and look longer at collections containing 4 or 16 items. Meanwhile, those infants who were habituated to 16 items were found to reliably dishabituate to collections containing 8 or 32 items. Since the infants in both conditions were found not to reliably dishabituate to collections containing (e.g.,) 12 items, these results were interpreted as showing that 6month-old infants can discriminate numbers, *provided they differ by a ratio of 1:2 or better*. But, on these grounds, Beck holds that their approximate number representations were not systematically recombining. For while the 6month-olds might have been capable of representing that *8 is less than (or different from) 16* and that *12 is less than (or different from) 24*, they were not able to represent that *12 is less than (or different from) 16*, since 12:16 is closer to 1:1 than 1:2.

To be clear, Beck's argument is not specific to the 6month-olds tested in Xu and Spelke's study. He formulates his argument in relation to approximate number representations posited in pigeons (Rilling & McDiarmid 1965) and his point is expected to apply no matter the acuity of the ANS in question. For instance, it would apply equally well to 9month-olds who have been said to reliably discriminate 2:3 ratios among numerical quantities (Lipton & Spelke 2003): here, a 9month-old who can represent 8 as smaller than 12 and 10 as smaller than 15 would (by hypothesis) not be capable of representing 10 as smaller than 12, *pace* systematicity. Additionally, Beck's argument does not turn on the precise numerical values represented. To simplify matters, I have been describing infants' approximate number representations as though they represent precise natural numbers (e.g., *8* or *8>6*). However, the same point would apply if infants were representing approximate numbers, like *8ish*, or precise numbers with a confidence estimation attached (Halberda 2016), "a blur on the number line" (Spelke & Tsivikn 2001) or some numerical range (like 7 to 9, or 7ish to 9ish – Ball 2017), etc. This is because systematicity would imply that the competence to represent 7–9 (or 8ish) is smaller than 14–18 (or 15ish), and the competence to represent 10–12 (or 11ish) is smaller than 20–24 (or 22ish), would imply the competence to represent 7–9 (or 8ish) is smaller than 10–12 (or 11ish) (at least insofar as '7-9' and '10-12', etc. feature as syntactically distinct elements in these representations, the structural relations among these elements remains constant, and the representations are composed from these elements).

So, does Beck's argument succeed in undermining the systematic recombinability, and by extension the compositionality, of these complex representations? It does not, and the failure of his critique should increase our confidence that approximate number representations genuinely compose.

The problem with Beck's argument is that it assumes a problematic interpretation of Weber's Law (a point noted by Halberda [2016] and acknowledged in Beck [forthcoming]). Specifically, it assumes what we might call *a cliff-edge-model of Weber discriminability*: namely that beyond some specific ratio, or threshold, discrimination becomes impossible. Thus, the argument, as it relates to Xu and Spelke's study (described above), assumes that when two numerical quantities differ by a ratio that is closer to 1:1 than some threshold (e.g., 1:2), subjects are no longer able to discriminate the quantities in question *at all* let alone represent one quantity as larger than the other. Performance simply falls off a cliff.

Admittedly, Xu and Spelke do suggest this, concluding from their study that number discriminations "*require* a 1:2 difference ratio" in the 6month-olds tested (Xu & Spelke 2000: B9 – my emphasis). Similar points apply to the specific studies on which Beck draws. Moreover, they are evident in seminal discussions of the ANS, as when Carey claims that (since adult humans require a 7:8 ratio or better to perform reliable discriminations) "The difference between eight and nine is not experienced at all, since eight and nine, like any higher successive numerical values, cannot be discriminated" (2009: 295). The point to stress is that, despite the prominence of these claims, the cliff-edge model is inaccurate.

For a start, the cliff-edge model conflicts with foundational work in psychophysics, where theories like Signal Detection Theory, jettison the idea of a *just noticeable perceptual difference* entirely (e.g., Lamming 1973). So, rather than performance falling off a cliff as some ratio threshold is exceeded, sophisticated and well-supported psychophysical theories predict that the discriminations facilitated by the ANS simply decrease *smoothly* in accuracy or reliability as ratios approximate 1:1. They, thus, predict no simple cut-off after which differences are not discriminated. With enough trials, one's ANS should (by hypothesis) be significantly above chance at discriminating 299 dots from 300 (or worse) provided the system can represent these quantities in the first place – e.g., when representing 299>149 and 300>150 (a made forcefully by Halberda and Odic [2014] and Halberda [2016]).

This sounds like a bold suggestion. But all available evidence supports it. In Barth et al.'s (2005) study (described in Section 2) we saw that 5year-olds discriminated 4:7 ratios more reliably than 4:6 ratios, and 4:6 ratios more reliably than 4:5 ratios. But at no (obvious) point did performance fall off a cliff. Similarly, while it is true that (e.g.,) Xu and Spelke (2000) report that 6month-old infants "require" a 1:2 difference for number discrimination to be possible, it has been noted that these researchers did not run enough trials for above chance discriminations of harder ratios to reveal themselves (Halberda 2016). The results, thus, fail to evince the cliff-edge model. Indeed, Libertus and Brannon (2010) replicated Xu and Spelke's findings and showed that as ratios became more favorable (than 1:2) performance increased smoothly, just as the abovementioned psychophysical models would predict.

Of course, these observations do not establish that there is *no* ratio threshold beyond which discriminations become *impossible*. However, a recent paper by Sanford and Halberda (forthcoming) speaks to this issue directly. They found that adult participants could perform statistically significant, above chance discriminations of 50:51 ratios, when sufficiently many trials were run. This is striking, not just because 50:51 is an extremely hard ratio (see Figure 3) – much harder than any ratio previously tested or described in studies of the ANS – but also because it was the hardest ratio these researchers examined, suggesting that harder ratios still would be discriminable also.

More importantly, Sanford and Halberda proceeded to compare participants' results against the predictions of two models: an idealized Signal Detection Theory model (which jettisons any notion of a just noticeable difference beyond some given ratio threshold) and a 'Give up model' "devised… in an attempt to capture the intuition that comparisons eventually become so difficult as to be imperceptible – i.e., that they give rise to random guessing" (9). For 398 of 410 participants tested, the former model accurately predicted their behavioral responses, and did so significantly better than the Give up model. So, while this empirical work is ongoing, a picture has emerged that is independently predicted by our best psychophysical modelling: rather than there being some just noticeable difference, or ratio threshold, after which ANS discriminations fall off a cliff, performance simply deteriorates gradually as ratios approximate 1:1. When ratios get difficult,

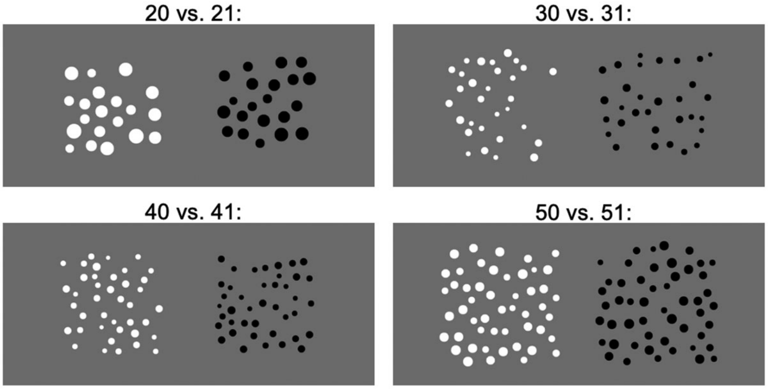

FIGURE 3 Illustration of ratio discriminations tested by Sanford and Halberda (forthcoming).

more trials must be run for above chance discrimination to reveal itself in the data. But genuinely *just noticeable differences*, of the sort assumed by the cliff-edge model, are little more than an intuitive fiction.

How does this bear on Beck's argument? It undermines it, because it suggests that approximate number representations *do* systematically recombine after all – just as the compositionality of these representations implies. Once we jettison the idea that there is some fixed ratio, beyond which numbers become indiscriminable, we seem forced to acknowledge that *any time* an ANS can represent two distinct quantities, or numbers, it possesses the competence and representational resources needed to (e.g.,) represent these as different and to represent one as larger than the other. So, if the ANS can represent 300 and 299 under any conditions (e.g., when representing 300>150 and 299>149), it possesses the representational resources needed to represent these as different (i.e., that 300≠299) and the former value as greater than the latter (i.e., that 300>299). What the system's conformity to Weber's Law means is simply that such discriminations will be highly error prone given 299:300's proximity to 1:1. As such, many trials will need to be run for significantly above chance performance to emerge in the data. But emerge it will, suggesting that errors ultimately pertain to performance limitations (for instance, noise in the encoding process), rather than speaking against an underlying competence to represent the differences and ordinal relationships in question – i.e., rather than speaking against that which the compositionality of these representations implies.

3.3 | Reason 3: Systematicity is best explained by a compositional structure

The fact that approximate number representations *seem* to feature in complex representations with a compositional structure coupled with the fact that *the* single most prominent argument against this suggestion fails, lends considerable support to the idea that approximate number representations *do* feature in compositionally structured representations. As a final remark, I wish to add that, in addition to making room for their compositionality, the systematic recombinability of these representations is positive evidence in favour of a compositional structure.

The reasons for this are familiar from seminal disputes in cognitive science, waged between classicists (who posit compositionally structured mental representations – e.g., Fodor & Pylyshyn 1988; Fodor & McLaughlin 1990) and connectionists (who deny this, positing non-compositionally structured representations with a distributed, context sensitive semantics – e.g., Smolensky 1988;

Chalmers 1990; Pollack 1990). Here, classicists point out that compositionally structured representations imply systematic competences in a system's operations, while non-compositional architectures do not. Hence, the *empirical* observation (Fodor 1994) that cognition exhibits systematic competences is said to be best explained by a classical architecture, or so the classicist contends.[4]

This is relevant to our discussion because we have now seen that approximate number representations systematically recombine. Somewhat surprisingly, an organism who can use their ANS to represent 50>25 and 51<100 possesses the competence to represent 50<51. Echoing the classicist, this is precisely what we would expect if these representations were compositionally structured. If representing 50>25 and 51<100 is (in either case) a function of representing the distinct elements these complex states comprise (e.g., 25, 50, 51 and 100 respectively), and the meaning of these complex states is determined by the meanings of these elements plus the way in which they are combined, then *any* creature with the competence and representational resources to represent 50>25 and 51<100 will possess the competence and representational resources needed to represent 50<51. This follows as a matter of "nomological necessity" (Fodor & McLaughlin 1990: 188).

By contrast, suppose that the elements in these representations have a context sensitive semantics, such that elements representing 50 and 51 in 50>25 and 51<100 have their semantics fixed by both their intrinsic properties and the contexts in which they are tokened (or perhaps just the latter). Here, it would not follow that the competence and representational resources needed to represent 50>25 and 51<100 implies the competence and representational resources needed to represent 50<51. Thus, something further would be needed to explain why systematicity obtains in the system, if and when it does, at which point the proponent of a compositional architecture can legitimately claim the upper hand. In much the way Darwin argued his evolutionary theory was to be favoured over creationism on account of its predicting and explaining observed similarities among cave forms in Europe and America, where creationism merely accommodated these (via ad hoc auxiliary posits), the proponent of compositionality can claim that their architecture is to be favoured over a non-compositional alternative (this example is borrowed from Aizawa 1997) on account of its predicting and explaining the observed systematicity under consideration (e.g., among approximate number representations) and not merely accommodating this via ad hoc auxiliary posits.

None of this is to deny that, in principle, performance limitations could still prevent a system employing compositionally structured representations from producing certain systematic recombinations of its representable elements (c.f. Aizawa 1997). Nor is it to deny that a non-compositional system *could* be wired up such that it exhibits the competences of a systematic cognizer (c.f. Chalmers 1990). The point is simply that compositional architectures straightforwardly predict, imply, and explain systematic *competences*, while non-compositional architectures do not. So, while these considerations may be non-decisive, the empirical fact that approximate number representations *do* systematically recombine is yet another mark in favour of thinking that they feature in compositionally structured complex representations.

4 | CONSTITUENCY AND ANALOGUE FORMAT IN THE ANS

The preceding discussion suggests that approximate number representations feature in complex compositionally structured states: they *seem to*, the most prominent argument to the contrary fails, and the empirical observation that these representations systematically recombine is most straightforwardly explained by the compositionality of the representations they comprise.

With these points in view, I'll now turn to the second (perhaps more controversial) aspect of their Privileged Compositionality: the suggestion that these compositionally structured representations possess a "canonical decomposition" into *privileged constituents*, despite possessing an *analogue* format.

To remind readers: Representations comprise privileged constituents when the semantics of a complex state depends on its decomposition into specific, semantically relevant, chunks or parts (rather than others). Linguistic sentences are the exemplars. As Fodor (2007) explains, the sentence 'John loves Mary' (understood to mean that *John loves Mary*) must be parsed [((John)$_{NP}$ ((loves)$_{V}$ (Mary)$_{NP}$) $_{VP}$)]. In so doing, chunks like 'John' and 'loves Mary' (and in turn, 'loves' and 'Mary') feature in its syntactic and semantic disambiguation and thus function as constituents of the sentence (they mean things which contribute to the sentence's overall meaning). Meanwhile, other sentence parts do not. For instance, 'John loves' and 'John… Mary' do not mean things which feature in the sentence's syntactic or semantic disambiguation and, thus, do not feature among the privileged constituents from which it is composed. Fodor and many others expect similar points to apply to the digital/linguiform representations posited in conceptual thought. I'm sympathetic to all the above.

What's critical is that theorists who embrace this conclusion often deny that privileged constituency obtains in complex analogue or iconic mental representations. For instance, Fodor (2007; 2008), Kosslyn (1980), Quilty-Dunn (2020), and Johnson-Laird (2001) all hold that these representations are akin to pictures in their conformity to *the picture principle*: just as you can (allegedly) "Take a picture of a person, [and] cut it into parts however you like" yet never be left with a part which fails to picture some part of the person or scene originally depicted, these researchers conclude that analogue mental representations "have interpretable parts, but they don't have constituents. Or, if you prefer, all the parts [of some such representation] is among its constituents" (Fodor: 2007: 108). So, while allowing that these representations are compositional, they deny that they are composed from *privileged* constituents in the above sense. Indeed, Carey (2009) seems to argue that this much applies to the approximate number representations we've been considering. For while she does not mention privileged constituency, she argues that these representations conform to the picture principle (135), holds that this much is implied (or perhaps entailed) by their non-digital, analogue format (*ibid.*), and claims that this is to be understood in precisely the same way as the abovementioned theorists (458) for whom a denial of privileged constituency looms large. I'll now argue that this is a mistake – approximate number representations *do* possess a canonical decomposition into privileged constituents *despite possessing an analogue format*, and the reasons for this help us to see that Fodor and Kosslyn were wrong to suggest otherwise in the case of visual imagery and early vision also.

4.1 | Why think approximate number representations analogue?

Characterizations of analogue representation vary in the literature. Analogue representations are often said to be 'continuous' in a way that digital representations are not (Goodman 1976), and we have just seen that some claim they are demarcated by their conformity to the picture principle, or an associated failure to canonically decompose into privileged constituents. However, another way of characterising analogue mental representations (which is not incompatible with these suggestions but is, in my view, more fruitful) appeals to an isomorphism, or homomorphism, between content-bearing properties of the vehicle and the magnitudes/intensities being represented (Lewis

1971; Maley 2011, forthcoming). As Von Neumann put it: "In an analogue machine each number is represented by a suitable physical quantity, whose values, measured in some pre-assigned unit, is equal to the number in question" (1958: 3). Thus, a mercury-in-glass thermometer is analogue, on this view, not because its mercury level varies continuously within the tube, but because it varies as a *linear* or (Von Neumann notwithstanding) a more broadly *monotonic* function of the temperature represented. This contrasts with the symbols of a digital thermometer which bear arbitrary similarities and differences to the numbers/temperatures they represent, like symbols in a language of thought.

While researchers studying the ANS are not always clear on this point, it is in this sense that their arguments support the claim that approximate number representations are analogue. These representations seem to represent numerical quantities by having some neuropsychological magnitude in the head vary as a monotonic function of these (Beck 2015; Carey 2009; Clarke 2022b). Thus, one magnitude (in the brain) is used to represent another magnitude (Peacocke 2019), by mirroring this (Beck 2015), and bearing some natural isomorphism (Burge 2018), or homomorphism (Maley forthcoming) towards it, irrespective of whether the vehicle is continuous (c.f. Gallistel & Gelman 2000) and irrespective of whether it conforms to the picture principle (c.f. Carey 2009).

Perhaps the most prominent argument in favour of this suggestion concerns the fact that an analogue format of this sort features in a simple and elegant explanation for the ANS's conformity to Weber's Law (Beck, 2015; 2019; Carey 2009; Condry & Spelke, 2008; Clarke 2022b; Dehaene 2011; Feigenson, Carey & Hauser, 2002; Feigenson, Dehaene & Spelke, 2004; Lipton & Spelke, 2003; Nieder & Miller 2004; Slaughter, Kamppi & Paynter, 2006). Consider a toy example:

Suppose that we've been tasked with keeping count of the goals scored at a local Scottish football match and decide to perform this task in an unorthodox manner. We take two identically proportioned buckets, assign one to each of the competing teams, and commit ourselves to pouring one glassful of water into the relevant bucket each time its corresponding team scores a goal. Proceeding in this way, we might reason that the bucket containing more water at the end of the match will belong to the winning team. Indeed, we might reason that the total volume of water that is in these buckets carries content about the number of goals each team has scored. This is because the number of goals scored by each team will be equal to the total volume of water in their corresponding bucket divided by the volume of a single glassful. Thus, if there is 1500ml of water in a bucket at the end of the game, and a single glassful contains 250ml, we might conclude that the corresponding team scored 6 goals. In this way, we can take the total volume of water to represent the number of goals scored by having the content-bearing magnitude (water volume) vary as a monotonic (or, more specifically, linear) function of the number of goals scored. That is, by serving as an *analogue* of the number represented, and representing it in an *analogue format* of the sort articulated above.

So far so good. But how does this relate to the ANS and its conformity to Weber's Law? The point to note is that the above example is idealized in at least one respect. It assumes that the amount of water in each glass is literally identical – e.g., that each glassful contains exactly 250ml of water. This is unlikely to be true of real-world systems. For instance, a glass used to fill one bucket is likely to be slightly bigger than a separate glass used to fill the other. Given the pace of local Scottish football matches, there may then be no time to ensure that each glassful contains *exactly* 250ml. But once this is acknowledged, noise could be expected to naturally accumulate in the buckets as more goals are counted. As a result, the accuracy or reliability with which the buckets' relative water volumes accurately track the winning team will end up a function of the ratio between the numbers of goals that each team scores. In other words, the accuracy with which these bucket representations pick out the winning team will be predicted by Weber's Law, just like the ANS's numerical discriminations.

To illustrate, suppose that while an average glassful contains 250ml of water, glassfuls vary in volume, such that one glass has a volume of 200-280ml and the other glass has a volume of 220-300ml. Here, we would still be able to tell who has won the match if the winner wins by two goals to one. After all, this would involve the winner's bucket containing somewhere between 400 and 600ml of water and, thus, more than the loser's bucket, which would only contain between 200 and 300ml of water. Indeed, we would be able to tell this much just as reliably any time the winner wins by scoring twice as many goals as their opponent. But, alas, matters will become less clearcut as the ratio in the number of goals scored by the winning and losing teams approximates 1:1. For instance, if the winning team wins by 4 goals to 3, there could be anywhere between 800 and 1200ml of water in their bucket and anywhere between 600 and 900ml in the loser's. Thus, there would be a range of possible scenarios in which the loser's bucket would contain more water than the winner's, with "errors" of this sort becoming ever more likely as the ratio in the number of goals scored approximates 1:1.

Of course, no one thinks that there are buckets of water inside our heads. What the example highlights is just that if the ANS represents numbers by having (functional or physical) magnitudes in the brain vary as a monotonic function of these, and we further acknowledge that noise could naturally accumulate in the abovementioned way, we can see why the ANS's numerical discriminations would conform to Weber's Law. In other words, the idea that analogue vehicles mirror the numbers they represent by varying in size as a monotonic function of these, features in a simple and elegant explanation for why the ANS might display its characteristic limitation. (As far as I can tell, this oft-repeated suggestion can be traced to Meck and Church's [1983] mode control model of counting and temporal representation in Norway rats.) None of the above turns on the continuity of the representations, however. Water volume is a (relatively?) continuous magnitude. But, for the purposes of the above example, we might as well pour cupfuls of discrete marbles into the buckets (Beck 2015). For if noise still accumulates as per above (and one cupful can contain anywhere between, e.g., 20 and 28 marbles, while the other can contain anywhere between 22 and 30 marbles) precisely the same points will apply: noise will accumulate such that the reliability with which the total number of marbles in each bucket picks out the winning/losing team ends up a function of the ratio of goals scored. Weber's Law still emerges in the system's discriminations.[5]

4.2 | From Analogue Format to the Picture Principle?

The preceding subsection introduces a widely endorsed reason for thinking that approximate number representations have an analogue format; for thinking that the vehicles of these representations represent the numerical quantities they do by varying as a monotonic function of these, and by thereby bearing some natural (and computationally efficacious) isomorphism towards these. Of course, moving from this to the conclusion that approximate number representations do, in fact, have an analogue format of this sort requires that there be no better proposals in the vicinity (for a sustained defence, see Beck 2019). But since the assumption that they do is commonplace in the literature and finds support from our best neuroscience (Nieder 2016), I'll assume it in what follows. The issue I would like to focus on is, instead, whether this threatens the idea that approximate number representations comprise privileged constituents in ways that Privileged Compositionality demands.

One reason why it might seem to threaten this, is that the above stated analogue format has been taken to imply (or perhaps entail) that approximate number representations conform to *the picture principle*, outlined at the start of this section. For instance, Carey (2009) endorses the abovementioned argument for their analogue format (although she uses line lengths, rather than volumes of water, to illustrate the suggestion) and takes this to imply that they are akin to realistic pictures in that parts of the representational vehicles represent parts of that which is represented by the representation as a whole. Since Carey is explicit that she intends this suggestion to be understood in the way that Kosslyn and others do when they characterise visual imagery (458), where a proposed rejection of privileged constituency looms large – i.e., where theorists claim that image parts can be "arbitrarily" defined (e.g., Kosslyn 1980: 30–33) with "every" part of the representation featuring in its compositional analysis (Fodor 2007: 107–8) – she is naturally read as suggesting that approximate number representations lack a canonical decomposition also: i.e., as claiming that *there is no privileged constituency in these representations*. In any case, seeing why this is (or would be) a mistake, enables us to see why it has been a mistake to characterize analogue mental representations in this way more generally – e.g., in the specific cases that Kosslyn and Fodor emphasise when framing the picture principle and explicitly rejecting the suggestion that analogue mental states decompose into privileged constituents.

To see why Carey thinks that approximate number representations conform to the picture principle, consider the buckets of water discussed in the previous section. These buckets were seen to represent numbers (of goals) by having content bearing properties of the vehicle (water volume) *increase* as a monotonic (linear) function of numbers represented. In particular: A bucket containing 2cupfuls represented *2* or *2 goals*, while a bucket containing 3cupfuls of water represented *3* (or *3 goals*), and so on. On these grounds, Carey notes that "the symbol for 3… contains the symbol for 2" as one of its proper parts (Carey 2009: 135), since 3cupfuls of water (representing 3) is composed of 2cupfuls of water (which represents 2 in and of itself) plus an additional third cupful of water (representing 1 in and of itself) (458). Since there is a natural sense in which 2 and 1 are parts of 3 (in a way that, e.g., 4 and 5 are not), Carey concludes that these analogue representations conform to the picture principle: parts of the representation represent parts of the numbers represented, in much the same way that parts of a photograph depict parts of the scene being pictured. Just "as a realistic picture of a dog representing a dog [has] parts of the symbol represent parts of the represented entity… Analog magnitude number representations are analog in this very sense" (*ibid*.).

Is Carey right to think this? An initial problem with her argument is that it turns on several dubious assumptions (for a full discussion, see Clarke 2022b). For one: Carey's argument seems to be founded on the (unargued) assumption that the content-bearing vehicle must *increase* in size with the number being represented. This was true in the bucket example described above. However, Weber's Law would emerge in the system's discriminations, and would emerge *for precisely the same reasons*, if the magnitude *decreased* in size as more goals were counted (ibid.; Beck 2015). For if we started with two *full* buckets of water and *removed* a cupful of water from the relevant bucket each time its corresponding team scored a goal, noise could now accumulate as more cupfuls are removed. As a result, Weber's Law could still emerge in the bucket system's discriminations (owing to the fact that the vehicles still mirror the numbers represented, in an analogue format). However, in this case, it would plainly be wrong to hold on to the suggestion that *any* parts of the vehicles (the bodies of water) are representing parts of the numbers being represented by these representations as a whole. If anything, a body of water representing *3goals*, might contain parts which represent 4, 5, and 6, but not parts which represent 1 or 2 – i.e., numerical quantities that are *not* parts of 3, rather than those that are, since water volume is inversely proportional to the numbers represented.

Of course, this merely highlights a *possibility* – it shows that (for all Carey has said) it is possible that parts of the analogue vehicles in question might fail to represent *any* parts of their content whatsoever. Nevertheless, this possibility is pertinent for our purposes insofar as conformity to the picture principle is meant to imply a lack of privileged constituency in analogue representations (Fodor 2007; Kosslyn 1980; Quilty-Dunn 2020). For in cases of the sort just outlined, our semantic analysis *must* depend on the magnitudes being parsed in privileged ways. Boldly stated: our interpretation of a given bucket representation will now require that the *whole* body of water (contained therein) be treated as *the* single semantically relevant representation of number. Thus, when the bucket representation represents that its corresponding team have scored 6 goals and, thus, less than another team, whose bucket indicates that they have scored 8 goals, mere portions of this body of water are not taken to mean things which feature in the semantic analysis of the complex state. As such, our semantic analysis of the representation depends on the content-bearing magnitude(s) having a canonical decomposition into one or more privileged constituents.[6]

Similar points also apply if we grant to Carey that parts of the neuropsychological magnitudes increase in size with the numbers being represented and, thus, possess parts which represent parts of the numerical quantities being represented by the representation as a whole. To see why, consider the original bucket system whose analogue vehicles increased in size when larger numbers were represented. Here, it may seem reasonable to follow Carey in supposing that the symbol for 3 contains a symbol for 2 and a symbol for 1 as its proper parts. However, the analogue representations of the bucket system still only represent whole *natural* numbers of goals. Hence, a slight increase/decrease in the volume of water in a bucket should not be taken to imply that (e.g.,) Tranent Juniors have scored 3.2 goals, let alone some irrational number of goals (e.g., π goals) – this is because, non-natural rational and irrational numbers (of goals) are not something the system functions to represent. So, what is happening when a bucket's water volume is taken to represent 7 goals? Presumably, it is being treated as though it comprises 7 distinct and privileged chunks, out of which its meaning is comprised. Meanwhile finer grained chunks of the water are not being taken to represent anything at all (e.g., a half chunk is not taken to represent 0.5 goals in and of itself). Perhaps these finer grained chunks could represent something under other circumstances (perhaps 9/10 of a chunk would be interpreted as a whole chunk had the water been parsed differently – i.e., had it possessed a different syntactic disambiguation); the point is just that, here, they are not, and they are thus failing to feature as constituents of the representation in *this* consumer's interpretation. Which is tantamount to saying that our interpretation of the bucket system's representations still involves a canonical decomposition of these representations into privileged constituents.

This latter point is worth considering since it is reasonable to suppose that the ANS primarily functions to represent whole or natural numbers, like the bucket system just described. For instance, in each of the studies described in Section 2, subjects were seen to compare and discriminate numbers of *whole items* (e.g., the numbers of *whole* dots in two displays). Indeed, well-known number illusions suggest that this much is not just typical, but integral to the functioning of the system. For instance, in the connectedness illusion, described previously, dots that are connected with thin lines (effectively turning pairs of dots into single dumbbell-shaped objects) cause collections in which they feature to appear significantly less numerous than otherwise identical arrays in which lines contain small breaks (Franconeri et al 2009) or in which the lines are free-floating (He et al. 2009). The reason for this seems to be that the ANS functions to enumerate whole bounded items (Dehaene & Changeux 1993; Franconeri et al. 2009; Spelke 1990; c.f., Green

2018), whose numerical quantities are most straightforwardly specified in terms of whole natural numbers (Clarke & Beck 2021a), raising the possibility that non-natural numbers lie beyond its expressive capacities.

Against this, it might be noted that there is provisional evidence that the ANS goes beyond representing natural numbers by representing fractions or ratios among these (ibid.). For instance, children are able to quickly identify the more favourable of two non-symbolic ratios (e.g., which of two arrays contains a more favourable ratio of blue to red gumballs), albeit imprecisely and in accord with Weber's Law – i.e., with the ratio *among* the ratios predicting the reliability with which subjects pick the more favourable of two collections (Szkudlarek & Brannon 2021). Hence, these ratio comparisons display a key signature limit of ANS processing. What's crucial to note is that, even here, the primacy of natural number representation is apparent. For instance, in a study that I have recently run with Chuyan Qu and Elizabeth Brannon (Qu et al. 2023), we've found that ratio comparisons are influenced by connectedness in the abovementioned way. In our study, subjects were (again) instructed to select the more favourable of two ratios (which of two gumball machines is most likely to produce a preferred gumball type), ignoring task irrelevant lines that populate the arrays. Despite these instructions, when preferred gumballs were visually connected with thin lines this reduced their perceived number and made the ratio in which they featured appear *less* favourable; conversely, connecting non-preferred gumballs reduced their perceived number and made the ratio in which they featured appear *more* favourable. This suggests that ratio (or rational number) extraction is still based on the prior processing of whole or natural numbers of preferred and unpreferred gumballs and that ratios are merely extracted and compared by subsequent mechanisms of approximate number representation. So, while future findings might upend these suggestions, extant evidence indicates that an initial stage of ANS processing produces approximate number representations of *natural number* akin to those of the bucket systems outlined above. When the system represents some relatively determinate quantity, like *7* or *6–8*, a neuropsychological magnitude varies as a monotonic function of this value. But even if the neuropsychological magnitude is, itself, continuous and increases in size in proportion to the numbers represented (points which, as we have seen, are themselves questionable), it is unlikely that every arbitrary part of the magnitude will represent some real number that is a part of that which gets represented by the magnitude as a whole. Rather, the magnitude will probably remain functionally chunky, breaking down into privileged chunks that represent whole natural numbers (i.e., function as constituents of the representation) leaving finer-grained parts of the magnitude semantically void and unable to feature in the ANS's syntactic decomposition.

4.3 | Why approximate number representations *must* comprise privileged constituents

The preceding discussion calls into question Carey's suggestion that approximate number representations conform to the picture principle. Moreover, it provides provisional reason to think that the individual representations of number that the ANS produces and manipulates possess a canonical decomposition into privileged constituents, irrespective of whether the analogue vehicles increase or decrease in size with the numbers being represented. As a final point, I wish to note that even if we put all this to one side (and grant that parts of the neuropsychological magnitudes implementing our approximate number representations simply possess arbitrarily defined parts which, themselves, always represent parts of the numbers represented by the vehicle as a whole) it *still* seems inescapable that these representations possess a canonical decomposition into privileged constituents when we consider the complex states they compose. So, while Carey's attempt to characterise approximate number representations in terms of the picture principle is best understood as an attempt to characterise individual numerical elements in these representations, an implied lack of canonical decomposition looks particularly untenable when we consider the complex representations they comprise.

To illustrate, recall that, in paradigm cases, approximate number representations do not simply represent things like *6!*. Rather, they seem to represent complex propositions, comprising syntactically distinct elements that pick out a collection and its numerical attribute respectively (Section 3). Thus, when I see a collection as 6 in number, my ANS might be more accurately said to represent the complex proposition:

([*that collection*]$_{\text{subject}}$ [*is 6 in number*]$_{\text{predicate}}$).

To compound matters, we saw that two (or more) numerical attributes can feature in a complex representation (and thus be compared) as when I consider two collections and come to represent these as containing 6 and 8 respectively and that 6<8 (i.e., by virtue of the two neuropsychological magnitudes implementing these numerical representations differing in size). Putting these claims together, we can then expect that in this example we would end up with a complex representation of something like the following sort:

([*that collection*]$_{\text{subject 1}}$ [*is 6 in number*]$_{\text{predicate 1}}$) < ([*that collection*]$_{\text{subject 2}}$ [*is 8 in number*]$_{\text{predicate 2}}$)

This is depicted below.

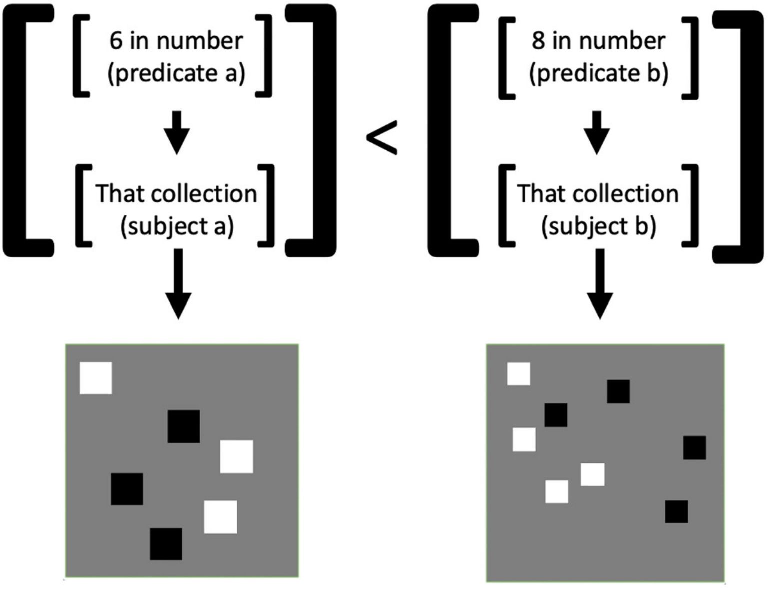

The point to note is that, here, there *must* be a canonical decomposition of the representation.

For a start, it seems vital to the content of this representation that predicate[1] is predicated of subject[1], and that predicate[2] is predicated of subject[2]. After all, these attributions are reflected in the content of the complex state. In this way, ([*that collection*]$_{\text{subject 1}}$ [*is 6 in number*]$_{\text{predicate 1}}$) and ([*that collection*]$_{\text{subject 2}}$ [*is 8 in number*]$_{\text{predicate 2}}$) mean things which contribute to the content of the complex state, suggesting that they might function as its constituents. And, indeed, these representation parts do seem to pass standard linguistic tests for constituency (e.g., tests listed by Sportiche et al. 2014: 48–57). For instance, we have already seen that these chunks can be systematically recombined with other chunks of the same (relevant) sort. Moreover, (when we bracket performance limitations) these chunks can be productively recombined, as in:

> ([*that collection*]$_{\text{subject 1}}$ [*is 6 in number*]$_{\text{predicate 1}}$)< ([*that collection*]$_{\text{subject 2}}$ [*is 8 in number*]$_{\text{predicate 2}}$) < ([*that collection*]$_{\text{subject 3}}$ [*is 10 in number*]$_{\text{predicate 3}}$)…

When they do, these complex chunks are treated like single primitive units in the representations in which they feature. For instance, they can be substituted for a single unstructured element without this changing the sentence's accuracy (and with this only causing a "little perturbation" in its meaning [ibid.: 50]). Thus, ([*that collection*]$_{\text{subject 1}}$ [*is 6 in number*]$_{\text{predicate 1}}$) might be replaced with a bare demonstrative '*that*' in the complex representation ([*that collection*]$_{\text{subject 1}}$ [*is 6 in number*]$_{\text{predicate 1}}$) < ([*that collection*]$_{\text{subject 2}}$ [*is 8 in number*]$_{\text{predicate 2}}$) to produce (that) < ([*that collection*]$_{\text{subject 2}}$ [*is 8 in number*]$_{\text{predicate 2}}$). For provided that the referent of the demonstrative is the relevant collection, this will not affect the accuracy of the representation.

For these reasons, complex chunks of *this* sort seem to function like constituents, despite comprising analogue elements. Yet when they do, other chunks or parts of the representation do not. In precisely the same way that 'John… Mary' is not a constituent of the sentence *John loves Mary* (while 'loves Mary' is), ([*that collection*]$_{\text{subject 1}}$ … [*is 8 in number*]$_{\text{predicate 2}}$) is *not* a semantically relevant part of the complex state described above. It does not mean *anything* (it has no accuracy conditions), and does not feature in our semantic analysis of the complex state described. For instance, there seems to be no sense in which this complex representation part can be manipulated, recombined, or analyzed as a single chunk (e.g., replaced with a single word or primitive element). And indeed, similar points even apply to the numerical elements themselves. For instance, when the system represents *6<8*, the values *6* and *8* feature as semantically relevant parts of the representation. But it is not true that the conjunction of any arbitrary part of the magnitude representing 6 and any arbitrary part of the magnitude representing 8 will mean something which features in the content of the representation (as a whole) even if we grant that arbitrary parts of the individual magnitudes would always represent real numbers in and of themselves. Thus, in each of these examples, it is simply not true that *every* arbitrary part of the representation represents a part of its represented content. Rather, these complex approximate number representations, comprising analogue parts, possess a canonical decomposition into certain *privileged constituents* from which their content is composed.

5 | PRIVILEGED COMPOSITIONALITY, GENERALIZED

I began this essay by asking whether analogue mental representations exhibit Privileged Compositionality, comprising compositionally structured *analogue* constituents. Focusing on the case of approximate number representation, I recommend an affirmative answer. These states *compose* into complex representations (they *seem* to, arguments to the contrary fail, and the systematic recombinability of these representations is best explained by their compositionality). And when they do, these representations are composed out of privileged constituents (this applies to both the complex representations in which they feature and the individual numerical magnitudes these comprise). Moreover, all of this obtains despite the fact that approximate number representations have an analogue format that distinguishes them from the digital representations emphasised in familiar formulations of Privileged Compositionality (e.g., the language of thought hypothesis).

To my mind, these points suggest that researchers tend to underestimate the syntactic complexity of analogue mental representations, as reflected in claims that analogue mental representations lack any notion of 'ill-formedness' (Haugeland 1981) or 'don't seem to have a syntax' (Kosslyn 1980). Nevertheless, it might be questioned whether these conclusions are in some way unique to the approximate number representations I have been discussing. After all, Fodor and Kosslyn's introduction of the Picture principle pertained to other analogue representations entirely – namely, "pre-attentive" perceptual representations and visual images, which might be thought more closely analogous to realistic pictures than the representations I've discussed. To close this essay, I wish to note that that all the above arguments apply equally well to these analogue representations also. The upshot is that Privileged Compositionality should probably be deemed a pervasive feature of analogue mental representation.

A full discussion of Fodor and Kosslyn's motivations for characterising early visual representation and visual imagery in terms of the Picture principle is beyond the scope of this article (see Clarke 2022a for more). But very briefly: Fodor thinks this much is motivated by the characteristically high storage capacity these states exhibit. For instance, while working memory only allows storage and processing of up to around four items at once (e.g., around four objects in the visual field – but see Fougnie & Alvarez 2011) early vision often seems to encode many items in parallel (e.g., Sperling 1960; Julesz 1961). According to Fodor this is best explained by the idea that these high-capacity early visual representations do not possess a canonical decomposition into privileged constituents representing individuals and their features respectively (Fodor 2007; 2008; see also: Quilty-Dunn 2020). Rather, they function more like uninterpreted photographs.[7] Arbitrarily defined parts of the representations simply depict corresponding regions of the visual scene, irrespective of whether those regions contain one item or many. So, when early vision represents a brown triangle, the brownness and triangularity of the representation gloop together "holistically", such that there is no decomposition into separate constituents or representational vehicles independently corresponding to the shape and color respectively (Quilty-Dunn 2020; see also: Dretske 1981; Davies 2021).

A challenge to this suggestion comes from decades of work on visual feature integration. In cases of *disjunctive* visual search, an 'oddball' will pop-out from a sea of distractor items when it is suitably distinguished along a single feature dimension. Thus, a single red circle will pop-out from a collection of green circles or red squares, and thus be identified as the odd one out, more or less as quickly when it is surrounded by 3 green circles/red squares or 30 (Treisman 1982). Indeed, similar points apply with many low-level visual properties (Wolfe & Horrowitz 2017) indicating that early ("pre-attentive") vision processes large portions of the visual field, efficiently and in parallel, to a sufficient level of granularity that items can be distinguished by *any one* of these features. But while this seems to be the kind of high-capacity visual process which motivates Fodor to suggest that *every* part of the early visual percept represents some portion of the seen scene, like a photograph, and *in a syntactically homogenous manner* (without breaking down into discrete constituents), other forms of visual search paint a more complex picture.

Most famously, in *conjunctive* search tasks, when a target item is defined by two features (e.g., by being the only red square, but not the only red item nor the only square item in the array) search times increase steadily with the number of distractors. Hence, it will (on average) take longer to find a target item in a sea of 30 distractors than in a sea of 29, and longer to find a target in a sea of 29 than 28. The traditional explanation for this has been that different feature types are represented in distinct 'feature maps', each of which makes salient oddballs that are defined by the specific feature type with which it is concerned; meanwhile, conjunctive search is slow because no single feature map picks out a conjunctively defined target – rather participants must engage in an effortful process, wherein overlapping locations on multiple maps are considered in series until oddballs are identified. Hence, the traditional explanation for the temporal asymmetry between disjunctive and conjunctive search tasks has involved the suggestion that distinct features enjoy an "independent psychological existence" in early vision (Treisman 1986).

What are we to make of this observation? An initial point to note is that this conflicts with the claims of those sympathetic to the Picture principle who maintain that early visual representations do not decompose into distinct parts or constituents that stand for distinct features, independently, and instead hold that these distinct features gloop together 'holistically' (Quilty-Dunn 2020; Davies 2021 – see: Clarke 2022a for discussion). That being said, it remains possible to hold onto the claim that these representations lack a canonical decomposition into *privileged* constituents. For while Fodor holds that *every* part of these representations represents some part of the visual field, he allows that this will only hold until the parts in question constitute some kind of representational primitive (like pixels) which no longer comprise semantically relevant parts (Fodor 2008: 173; Quilty-Dunn 2020: fn.8). Thus, he might respond that even if these psychological primitives (or concatenations thereof) do not encode multiple properties (like colour and shape) holistically, there remains a sense in which any given *primitive element* or *collection of primitive elements* represents some part of the visual scene. So, while 'John… Mary' is not a constituent of 'John loves Mary', a Fodorian might maintain that Frankensteinian percept parts, e.g., comprising a primitive representation of blue in the bottom left of my visual field and primitive representation of vertically-oriented-edgeness in the top-right, will always represent some part of what I see, holding that this is true for any arbitrary collection of primitives such that no complex part is privileged over any other portion of the early visual representation.

This is where our discussion of the ANS becomes relevant: To see why this won't do, consider the fact that while the high-capacity representations of early vision encode distinct feature types independently, early vision is capable of integrating features across large portions of the visual field, and not simply at the locus of focal attention. Indeed, this much was evident in the connectedness illusion, described several times previously in connection to the ANS. Here, the visual system seems to enumerate (and thus process) whole bounded objects across the visual field – hence why connecting items with thin dots (and turning them into single dumbbell shaped ovjects) reduces their perceived number, even when subjects try to ignore the lines completely (Franconerri et al. 2009; He et al. 2009; Kirjakovski & Matsumoto 2016). But notice: once it is acknowledged that the visual system integrates distinct features into bounded objects across large portions of the visual field (and outside of focal attention) it seems that these complex representations *must* possess a canonical decomposition into privileged constituents, and for precisely the same reasons as the complex approximate number representations I focussed on in this essay. In precisely that same way that 'John… Mary' is not a constituent of the sentence *John loves Mary* (while 'loves Mary' is), and ([*that collection*]$_{\text{subject 1}}$ … [*is 8 in number*]$_{\text{predicate 2}}$) is *not* a semantically relevant part of the complex ANS representation ([*that collection*]$_{\text{subject 1}}$ [*is 6 in number*]$_{\text{predicate 1}}$)< ([*that collection*]$_{\text{subject 2}}$ [*is 8 in number*]$_{\text{predicate 2}}$), ([brown]…[splodge]) is not a constituent of a complex visual state which attributes *brown* to *triangle* in the bottom left of the visual field and *grey* to *splodge* in the top right. Thus, it seems that – like approximate number representations – these complex states possess a canonical decomposition into privileged constituents from which they compose. They are not simply syntactically unstructured concatenations of primitive elements (*pace* Kosslyn 1980: 30–33; *pace* Kosslyn et al. 2006: 13; *pace* Fodor 2007; 2008; *pace* Quilty-Dunn 2020). (For a fantastic discussion of related issues, see: Lande 2020).

If these observations suggest that high-capacity early visual representations comprise privileged constituents, and do not conform to the Picture principle as Fodor and Kosslyn frame it, why think them analogue? Once again, our discussion of the ANS is instructive. These early visual states appear to be analogue in much the same way as approximate number representations – they represent magnitudes and intensities by having some content-bearing magnitude in the head vary as a monotonic function of the values represented. One reason to say this is that the discriminability of early visual representations (of the sort under consideration) also conforms to Weber's Law. For instance, in cases of disjunctive visual search, a long line will pop out from an array of short lines. However, the strength or reliability with which the long line pops out is predicted by the relative ratio of the lines' lengths – thus the line will pop out more reliably if it differs from distractor lines by a ratio of 4:3 than 8:7 (Treisman 1986). Crucially, we have seen that while this does not threaten the compositionality and systematic recombinability of these representations (c.f. Beck 2012), it is evidence that analogue constituents are involved (Beck 2019). Indeed, converging evidence for the postulation of analogue feature representations in early vision is not hard to come by. For one thing, this explains various (otherwise puzzling) asymmetries in disjunctive visual search – for instance, why it is that a long line will pop out from an array of short lines, but a short line is significantly less likely to pop out from an array of long-lines (Treisman 1986). On the account recommended here, this is explained by the analogue format of the representations involved – it is the visuo-computational equivalent of finding the longest spaghetti in the pack by banging the butt of the packet on the table and seeing which piece sticks out furthest (Clarke 2022a). In any case, the postulation of these analogue states is fully consistent with the Privileged Compositionality these analogue representations seem to possess.

6 | CONCLUSION

Analog mental representations are often treated as syntactically unstructured representations, lacking in Privileged Compositionality – as representations which fail to systematically recombine (Beck 2012), lack a canonical decomposition into privileged constituents (Fodor 2007; Kosslyn 1980) and/or defy any notion of 'ill-formedness' (Haugeland 1981; Kosslyn 1980). The present treatment has argued that all of this is a mistake. Focussing on the ANS, I proposed that its representations feature in complex and compositionally structured representations, which comprise privileged constituents, despite possessing a distinctively non-linguistic analogue format wherein neuropsychological magnitudes represent numbers by varying as a monotonic function of these. In addition, I have suggested that these states provide a model for understanding other putatively analogue representations in the human mind. For instance, I have suggested that while the states of early vision do not conform to the picture principle in the way that Fodor and others maintain, there is considerable reason to think them analogue in much the way our approximate number representations are, with neuropsychological vehicles representing visual magnitudes and intensities by varying a monotonic function of these (as with the ANS, this explains why their discriminability would conform to Weber's Law; in addition, it can explain the presence of various asymmetries in visual search – e.g., why large/bright/oblique items pop-out from a sea of small/dull/vertical items but not vice versa). And since these distinct analogue elements are structured and bound together in highly specific ways across the visual field (e.g., as illustrated in the connectedness illusion), it is not plausible to deny these states a canonical decomposition into privileged constituents from which they comprise. Rather, analogue features are bound together in syntactically complex and specific ways across the visual field. A complete understanding of these analogue, iconic, or otherwise picture-like states will, thus, require an appreciation for their constituent structure and grammatical principles by which they compose.

ACKNOWLEDGEMENTS

Thanks to Liz Brannon and Jake Beck for helpful and constructive comments on an earlier version of this paper, and to Brett Karlan and Ben Henke for discussion of these issues. Thanks also to David Barack for detailed comments – I think David rejects every conclusion drawn in this paper which (sincerely!) helped convince me that they are worth defendin

ENDSNOTES

[1] A further complication concerns whether the abovementioned results are better understood to involve a representation of *cardinality* as opposed to *number*, per se (Opfer et al. 2021; but see Clarke & Beck 2021b). This is an interesting possibility, deserving of discussion. However, nothing I'll say here turns on this issue.

[2] Beck (2015) makes a related point. But in addition to maintaining some syntactic independence among a representation of the collection and its numerical attribute, he holds that there is a further independence in the representation of the *magnitude* and the *magnitude type* being represented. Thus, he suggests that there is a component of the representation which says *this big* and an independent component which tags this magnitude as a *number* as opposed to (say) a *distance*. Beck endorsed this suggestion because he was impressed by evidence for a generalized magnitude system, according to which magnitude representations (quite generally) enjoy some kind of neuropsychological overlap. I'm impressed by evidence *against* the existence of a generalized magnitude system (Odic 2018; Cassanato & Pitt 2019). But nothing I say here hangs on this. If I'm wrong to deny the existence of generalized magnitudes, processed by a generalized magnitude system, then I would suggest that there is simply *more* compositional complexity to these representations than that which is described here.

[3] Interestingly, Beck does not take this breakdown of systematicity to refute the compositionality of approximate number representations. Indeed, Beck maintains that these representations can "be conceptualized as an ordered triplet", *composing* from an element which represents the size of the magnitude, a separate element which specifies the magnitude type (e.g., number as opposed to distance), and a third element which specifies the object or collection to which the magnitude is attributed (2012: 592). Intuitively, however, this is because these elements are (by hypothesis) capable of systematic recombination. Thus, if one can represent *collection[1]* as *6 in number* and *collection[2]* as *8 in number*, Beck should maintain that one possesses the competence and representational resources to represent *collection[2]* as *6 in number*, owing to the compositionality of these states. The trouble is, this seems to be in tension with the claim that distinct numerical attributes fail to systematically recombine (at least insofar as distinct approximate number representations are supposed to feature as semantically relevant parts of complex compositional states, as Beck seems to maintain). For if the ability to represent one collection as *X in number* (as such) is dependent on the number of elements attributed to other collections, then it is not true that individual numerical elements make fixed semantic contributions to every (complex) approximate number representation in which they feature, and the complex states seem not to be genuinely compositional.

[4] Clark (1990) objects to this suggestion on behalf of the connectionist. He maintains that the systematicity of thought is not an empirical fact at all, in need of some deep mechanistic explanation (*pace* Fodor). It is, instead, a conceptual fact that when we encounter a creature who (otherwise) seems able to represent *aRb*, the discovery that they are fundamentally incapable of representing *bRa* undermines the grounds we had for crediting them with genuine *aRb* representation to begin with. The fact that Beck (2012) questions the systematic recombinability of representable elements in approximate number representation, without taking this to undermine the ANS's capacity to represent (somewhat determinate) numerical quantities (in certain contexts) suggests to me that this is a mistake – it is not true that we are unwilling to credit creatures with aRb thoughts *unless* we are also willing to credit them with bRa thoughts. Moreover, the fact that Beck's argument turned on mistaken *empirical* assumptions, suggests that the systematic recombinability of *these representations* really is an *empirical* discovery – it is a fact that cries out for some form of mechanistic explanation.

[5] In both examples, water volume/number of marbles increases as a *linear* function of number, and noise accumulates by some (relatively) stable amount with each goal counted. Weber's Law could also emerge if these magnitudes varied as a *logarithmic* function of represented number, with differences in water volume/number of marbles becoming ever more compressed as more goals are counted, provided that noise levels do not accumulate in the above way but vary from an ideal level across some standard deviation. Indeed, researchers characterising approximate number representation often state that content-bearing properties of the neuropsychological vehicles vary as a logarithmic function of represented numbers in this way. My own view is that the linear model is better placed to capture the ANS's computations. For instance, the linear model offers a simple explanation for how the ANS might perform the addition operations described in Section 2. On this view, the ANS might add two quantities together by simply adding together the vehicles used to represent either collections' quantity – this is analogous to the way that, at the end of the football match, we might pour one bucket of water into the other and take the total volume of water in the latter bucket to carry content about the total number of goals scored by *both* teams. This wouldn't work if the mapping was logarithmic, rendering the addition operation more complex. While it might be replied (as Jake Beck did, in comments on an earlier draft) that a logarithmic model could simplify other arithmetic operations, like multiplication or division of one approximate number representation by another by reducing these to addition/subtraction operations (e.g., Log[a] x Log[b] = Log[a+b]; Log[a]/Log[b] = Log[a-b]), it is questionable whether the ANS facilitates such operations. For instance, while ANS multiplication has been observed, a recent study by Pickering*et al.* (2022) tested subjects' abilities to multiply ANS representations by numbers between 2 and 8 and found that while subjects were above chance in multiplying approximate quantities by 2-4 (i.e., by numbers within the subitizing range) performance fell to chance with numbers 5-8 (outside of the subitizing range and [by hypothesis] in need of representation by the ANS). This led the authors to conclude that "two ANS representations cannot be multiplied together" (1). In any case, nothing I say turns on this. Whether the mapping is linear or logarithmic, the vehicles involved represent numbers by *varying as a monotonic function of these*: they thereby qualify as analogue in the way I am using the term.

[6] Some theorists might have expected otherwise because they deem the system's conformity to Weber's Law evidence that the vehicles employed are *continuous* (e.g., Dehaene 2011; Gallistel & Gelman 2000). After all, they seem to be in the bucket example (between any two values or water volumes there could always seem to be a third). This is suggestive because when theorists claim of analogue representations that

"Variations are smooth or continuous, without "gaps"",

they are prone to suggesting that

unlike switches, abacuses, or alphanumeric inscriptions, every (relevant) setting or shape is allowed – *nothing is ill-formed* (Haugeland 1981: 220)

and that

These representations don't seem to have a syntax: (Kosslyn 1980: 30).

But even bracketing the point (noted above) that the ANS's conformity to Weber's Law does not imply continuous vehicles, the present example highlights the need to distinguish two issues: even if we grant to Haugeland that any given volume of water in the analogue bucket representation *would* count as a well-formed representation for the system (and there really is no notion of ill-formedness here – which I doubt, for reasons that will become clear), it does not follow that every part of a continuous analogue vehicle features as a constituent of the representation it comprises (i.e., meaning something which features in its compositional analysis – c.f. Kosslyn 1980; Fodor 2007). In the current example, the body of water (taken as a whole) features as a constituent of the representation, but portions of that body of water do not. Thus, privileged constituency is compatible with analogue vehicles, even when they are fully continuous.

[7] Fodor actually seems to oscillate on this point. While the idea that these representations are left (largely) uninterpreted is central to his oft-repeated claim that analogue states fail to *represent as*, he occasionally seems to hold that the high-capacity representations of early vision drive high-capacity visual computations (e.g., binocular integration in random-dot stereograms – Julesz 1971; see: Fodor 2008), with others appealing to the high-capacity representations of early vision to explain (e.g.,) how ensemble percepts are extracted efficiently from large collections (Quilty-Dunn 2020). In either case, this suggests that the high-capacity representations of early vision have relatively determinate contents (i.e., they are not *uninterpreted* and wildly ambiguous), such that they can inform high-capacity, parallel computations.

REFERENCES

Aizawa, K. (1997), Explaining Systematicity. *Mind & Language*, 12: 115–136. https://doi.org/10.1111/j.1468-0017.1997.tb00065.x

Ball, B. (2017). On representational content and format in core numerical cognition. *Philosophical Psychology*, 30(1–2), 119–139. https://doi.org/10.1080/09515089.2016.1263988

Barth, H., Kanwisher, N., & Spelke, E. (2003). The construction of large number representations in adults. *Cognition*, 86(3), 201–221. https://doi.org/10.1016/S0010-0277(02)00178-6

Barth, H., La Mont, K., Lipton, J., & Spelke, E. S. (2005). Abstract number and arithmetic in preschool children. *Proceedings of the National Academy of Sciences*, 102(39), 14116–14121. https://doi.org/10.1073/pnas.0505512102

Barth, H., La Mont, K., Lipton, J., Dehaene, S., Kanwisher, N., & Spelke, E. (2006). Non-symbolic arithmetic in adults and young children. *Cognition*, 98(3), 199–222. https://doi.org/10.1016/j.cognition.2004.09.011

Beck, J. (2012) The Generality Constraint and the Structure of Thought. *Mind*, 121, 483: 563–600, https://doi.org/10.1093/mind/fzs077

Beck, J. (2015). Analogue Magnitude Representations: A Philosophical Introduction. *The British Journal for the Philosophy of Science*, 66(4), 829–855. https://doi.org/10.1093/bjps/axu014

Beck, J. (2019). Perception is Analog: The Argument from Weber's Law. *The Journal of Philosophy*, 116(6), 319–349. https://doi.org/10.5840/jphil2019116621

Beck, J. (forthcoming). Contents and Vehiles in Analog Perception. *Crítica*.

Burge, T. 2018. *Iconic Representation: Maps, Pictures and Perception*. In: Wuppuluri S., Doria F. (eds) The Map and the Territory. The Frontiers Collection. Springer, Cham.

Carey, S. (2009). *The origin of concepts*. Oxford: Oxford University Press.

Chalmers, D. (1990). Syntactic Transformations on Distributed Representations. *Connection Science*, 2(1-2): 53–62.

Clark, A. (1990) *Microcognition: Philosophy, Cognitive Science, and Parallel Distributed Processing*. MIT Press.

Clarke, S. (2021). Cognitive Penetration and Informational Encapsulation: Have we been failing the module? *Philosophical Studies*. https://doi.org/10.1007/s11098-020-01565-1

Clarke, S. (2022a). Mapping the Visual Icon. *The Philosophical Quarterly*, 72, 3: 552–577, https://doi.org/10.1093/pq/pqab054

Clarke, S. (2022b). Beyond the icon: Core cognition and the bounds of perception. *Mind & Language*. https://doi.org/10.1111/mila.12315

Clarke, S. & Beck, J. (2021a). The number sense represents (rational) numbers. *Behavioral and Brain Sciences*, 44, E178. https://doi.org/10.1017/S0140525X21000571

Clarke, S. & Beck, J. (2021b). Numbers, numerosities, and new directions. *Behavioral and Brain Sciences*, 44, E205. https://doi.org/10.1017/S0140525X21001503

Condry, K. F., & Spelke, E. S. (2008). The development of language and abstract concepts: The case of natural number. *Journal of Experimental Psychology: General*, 137(1), 22–38. https://doi.org/10.1037/0096-3445.137.1.22

Cordes, S., Gelman, R., Gallistel, C. R., & Whalen, J. (2001). Variability signatures distinguish verbal from nonverbal counting for both large and small numbers. *Psychonomic Bulletin & Review*, 8(4), 698–707.

Davies, W. (2021). Colour Relations in Form. *Philosophy and Phenomenological Research*, 102: 574–594. https://doi.org/10.1111/phpr.12679

Dehaene, S. (2011). *The Number Sense: How the Mind Creates Mathematics, Revised and Updated Edition: How the Mind Creates Mathematics*. Oxford University Press, USA.

Dehaene S, Changeux JP. (1993). Development of elementary numerical abilities: a neuronal model. *J Cogn Neurosci*. 1993;5(4):390-407. https://doi.org/10.1162/jocn.1993.5.4.390

DeWind, N.K., Bonner, M.F., & Brannon, E. (2020). Similarly oriented objects appear more numerous. *Journal of Vision* 20(4):4. doi: https://doi.org/10.1167/jov.20.4.4.

Di Lollo, V. (2012). The feature-binding problem is an ill-posed problem. *Trends in Cognitive Sciences*, 16(6), 317–321. https://doi.org/10.1016/j.tics.2012.04.007Dretske 1980;

Dretske, F. (1981). *Knowledge and the Flow of Information*. Cambridge: MIT Press.

Echeverri, S. (2017). Visual Reference and Iconic Content. *Philosophy of Science* 84 (4):761-781.

Feigenson, L. (2011). Objects, sets, and ensembles. In S. Dehaene & E. Brannon (Eds.), *Space, time and number in the brain: Searching for the foundations of mathematical thought* (pp. 13–22). Elsevier Academic Press. https://doi.org/10.1016/B978-0-12-385948-8.00002-5

Feigenson L, Carey S, Hauser M. The representations underlying infants' choice of more: object files versus analog magnitudes. *Psychol Sci*. 2002 Mar;13(2):150-156. https://doi.org/10.1111/1467-9280.00427. PMID: 11933999.

Feigenson L, Dehaene S, Spelke E. Core systems of number. *Trends Cogn Sci*. 2004 Jul;8(7):307-314. https://doi.org/10.1016/j.tics.2004.05.002. PMID: 15242690.

Fodor, J. A. (1994). *The elm and the expert: Mentalese and its semantics*. The MIT Press.

Fodor, J. 2007. Revenge of the Given. In Brian P. McLaughlin & Jonathan D. Cohen (eds.), *Contemporary Debates in Philosophy of Mind*. Blackwell. pp. 105–116

Fodor, J. 2008. *LOT2*. Cambridge: MIT Press.

Fodor & Pylyshyn, 1988. Connectionism and cognitive architecture: A critical analysis. *Cognition*, 28(1-2): 3–71.

Fodor, J., & McLaughlin, B. P. (1990). Connectionism and the problem of systematicity: Why Smolensky's solution doesn't work. *Cognition*, 35(2), 183–204. https://doi.org/10.1016/0010-0277(90)90014-B

Fougnie, D., & Alvarez, G. A. (2011). Object features fail independently in visual working memory: Evidence for a probabilistic feature-store model. *Journal of Vision*, 11(12), Article 3. https://doi.org/10.1167/11.12.3

Franconeri, S. L., Bemis, D. K., & Alvarez, G. A. (2009). Number estimation relies on a set of segmented objects. *Cognition*, 113(1), 1–13. https://doi.org/10.1016/j.cognition.2009.07.002

Gallistel, R., & Gelman, R. (2000). Non-verbal numerical cognition: From reals to integers. *Trends in Cognitive Sciences*, 4(2), 59–65.

Gebuis, T., Cohen Kadosh, R., & Gevers, W. (2016). Sensory-integration system rather than approximate number system underlies numerosity processing: A critical review. *Acta Psychologica*, 171, 17–35. https://doi.org/10.1016/j.actpsy.2016.09.003

Goodman, N. (1976). *Languages of Art: An Approach to a Theory of Symbols*, 2nd edition, Indianapolis: Hackett Publishing Company.

Green, E.J. (2018) What do object files pick out? *Philosophy of Science*, 85(2): 177–200.

Green, EJ. & Quilty-Dunn, 2020. What is an Object File? *British Journal for the Philosophy of Science*. axx055.

Greenberg, G. (2013). Beyond Resemblance. *Philosophical Review* 122 (2):215-287.

Halberda, J. (2016). Epistemic Limitations and Precise Estimates in Analog Magnitude Representation. In *Oxford series in cognitive development. Core knowledge and conceptual change* (pp. 171–190).

Halberda, J., & Odic, D. (2014). The precision and internal confidence of our approximate number thoughts. In D. C. Geary, D. Berch, & K. Koepke (Eds.), *Evolutionary origins and early development of number processing* (pp. 305–335). Cambridge, MA: Academic Press.

Halberda, J., Sires, S. F., & Feigenson, L. (2006). Multiple Spatially Overlapping Sets Can Be Enumerated in Parallel. *Psychological Science*, 17(7), 572–576. https://doi.org/10.1111/j.1467-9280.2006.01746.x

Haugeland, J. (1981). Analog and Analog. *Philosophical Topics*, 12(1), 213–225. http://www.jstor.org/stable/43153852

He, L., Zhang, J., Zhou, T., & Chen, L. (2009). Connectedness affects dot numerosity judgment: Implications for configural processing. *Psychonomic Bulletin & Review*, 16(3), 509–517. https://doi.org/10.3758/PBR.16.3.509

Izard, V., Sann, C., Spelke, E. S., & Streri, A. (2009). Newborn infants perceive abstractnumbers. *Proceedings of the National Academy of Sciences*, 106(25), 10382–10385. https://doi.org/10.1073/pnas.0812142106

Johnson-Laird, P. N. (2001). Mental models and human reasoning. In E. Dupoux (Ed.), *Language, brain, and cognitive development: Essays in honor of Jacques Mehler* (pp. 85–102). The MIT Press.

Josh, J., Feigenson, L., & Halberda, J. (2007). Infants ability to enumerate multiple spatially-overlapping sets in parallel. *Journal of Vision*, 7(9): 220. https://doi.org/10.1167/7.9.220

Julesz, B. (1961). *Principles of Cyclopean Vision*. MIT Press.

Khooshabeh and Hegarty, 2008. Differential effects of color on mental rotation as a function of spatial ability. In: International Spatial Cognition Conference. Freiburg, Germany

Kirjakovski, A., & Matsumoto, E. (2016). Numerosity underestimation in sets with illusory contours. *Vision Research*, 122, 34–42. https://doi.org/10.1016/j.visres.2016.03.005

Kosslyn, S. 1980. *Image and Mind*. Cambridge: Harvard University Press.

Lamming, D. (1973). *Mathematical Psychology*. New York: Academic Press.

Kosslyn, S., Thompson, W. & Ganis, G. (2006). *The Case for Mental Imagery*. Oxford: OUP.

Lande, K. J. (2020). Mental structures. *Noûs*, 55, 649–677. https://doi.org/10.1111/nous.12324

Lande, K. (2021). Mental Structures." *Noûs* 55(3): 649–677.

Leibovich, T., Katzin, N., Harel, M., & Henik, A. (2017). From 'sense of number' to 'sense of magnitude' – The role of continuous magnitudes in numerical cognition. *Behavioral and Brain Sciences*, 1–62. https://doi.org/10.1017/S0140525X16000960

Lewis, D. (1971). Analog and digital. *Noûs*, 5(3), 321–327. https://doi.org/10.2307/2214671

Libertus ME, Brannon EM. (2010). Stable individual differences in number discrimination in infancy. *Developmental Science*. 13(6):900-906. https://doi.org/10.1111/j.1467-7687.2009.00948.x.

Lipton, J.S. & Spelke, E.S. (2003). Origins of Number Sense: Large-Number Discrimination in Human Infants. *Psychological Science*, 14(5), 396–401.

Lyons, J. (2021). Contents of the Approximate Number System. *Behavioral and Brain Sciences*, 44.

Maley, C. J. (2011). Analog and digital, continuous and discrete. *Philosophical Studies*, 155(1), 117–131. https://doi.org/10.1007/s11098-010-9562-8

Maley, C.J. (forthcoming). Icons, magnitudes, and their parts. *Crítica*.

Mandelli, M.J.F., & Kiper, D.C. (2005). The local and global processing of chromatic Glass patterns. *Journal of Vision*, 5(5):2. doi: https://doi.org/10.1167/5.5.2.

McCrink, K., & Spelke, E. S. (2010). Core multiplication in childhood. *Cognition*, 116(2), 204–216. https://doi.org/10.1016/j.cognition.2010.05.003

McCrink, K., & Spelke, E. S. (2016). Non-symbolic division in childhood. *Journal of Experimental Child Psychology*, 142, 66–82. https://doi.org/10.1016/j.jecp.2015.09.015

McCrink, K., & Wynn, K. (2004). Large-Number Addition and Subtraction by 9-Month-Old Infants. *Psychological Science*, 15(11), 776–781. https://doi.org/10.1111/j.0956-7976.2004.00755.x

Meck, W. H., & Church, R. M. (1983). A mode control model of counting and timing processes. *Journal of Experimental Psychology: Animal Behavior Processes*, 9(3), 320–334. https://doi.org/10.1037/0097-7403.9.3.320

Neisser, U. (1967). *Cognitive Psychology*. Routledge.

Nieder, A. (2016). The neuronal code for number. *Nature Reviews Neuroscience*, 17(6), 366–382. https://doi.org/10.1038/nrn.2016.40

Nieder, A. & Miller, E. K. (2004). A parieto-frontal network for visual numerical information in the monkey. *Procedings of the National Academy of Sciences*, 101(19):7457-7462. https://doi.org/10.1073/pnas.0402239101.

Odic, D., & Halberda, J. (2015) Eye movements reveal distinct encoding patterns of number and cumulative surface area in random dot arrays. *Journal of Vision*,15, 15-15.

Odic, D. (2018). Children's intuitive sense of number develops independently of their perception of area, density, length, and time. *Developmental Science*, 21(2), e12533. https://doi.org/10.1111/desc.12533

Opfer, J., Samuels, R., Shapiro, S., & Snyder, E. (2021). Unwarranted philosophical assumptions in research on ANS. *Behavioral and Brain Sciences* 44.

Peacocke, C. (2019). *The Primacy of Metaphysics*. Oxford: OUP.

Pickering, J., Adelman, J., & Inglis, M. (2022). Are Approximate Number Representations Numerical? Journal of Numerical Cognition.

Pollack, Jordan B. (1990). Recursive distributed representations. *Artificial Intelligence* 46 (1-2):77-105.

Prinzmetal, W. 2012. At the core of feature integration theory: on Treisman and Schmidt (1982). In J. Wolfe & L. Robertson (Eds.), Oxford series in visual cognition. From perception to consciousness: Searching with Anne Treisman (p. 211–216).

Pylyshyn, Z. 2007. Multiple Object Tracking. *Scholarpedia*, 2(10):3326.

Qu, C., Szkudlarek, E., & Brannon, E.M. (2021). Approximate multiplication in young children prior to multiplication instruction. *Journal of Experimental Child Psychology*, 207, 105116. https://doi.org/10.1016/j.jecp.2021.105116

Qu, C., DeWind, N., & Brannon, E. (2022). Increasing entropy reduces perceived numerosity throughout the lifespan. *Cognition*, 225, 105096. https://doi.org/10.1016/j.cognition.2022.105096

Qu, C., Clarke, S., & Brannon, E. M. (2023). Rational Number Representation by the Approximate Number System. Available at SSRN: https://ssrn.com/abstract=4486602 or http://dx.doi.org/10.2139/ssrn.4486602

Quilty-Dunn, J. (2020) Perceptual Pluralism. *Noûs*. 54(4): 807–838.

Rilling, M., & McDiarmid, C. (1965) 'Signal Detection in Fixed Ratio Schedules'. *Science*, 148, 526–527.

Sanford, E. and Halberda, J. (forthcoming). Successful discrimination of tiny numerical differences. *Journal of Numerical Cognition*.

Shepard, R. N. (1978). The mental image. *American Psychologist*, 33(2), 125–137. https://doi.org/10.1037/0003-066X.33.2.125

Slaughter V, Kamppi D, Paynter J. (2006). Toddler subtraction with large sets: further evidence for an analog-magnitude representation of number. *Developmental Science*. 9(1):33-39. https://doi.org/10.1111/j.1467-7687.2005.00460.x.

Smolensky, P. (1988). On the proper treatment of connectionism. *Behavioral and Brain Sciences*, 11(1), 1–74. https://doi.org/10.1017/S0140525X00052432

Spelke, E. S. (1990). Principles of object perception. *Cognitive Science*, 14(1), 29–56

Spelke, E., & Tsivkin, S. (2001). Language and number: a bilingual training study. *Cognition*,78, 45–88.

Sperling, G. 1960. The information available in brief visual presentations. *Psychological Monographs: General and Applied*, 74(11): 1–29.

Sportiche, D., Koopman, H., & Stabler, E. (2014). *An Introduction to Syntactic Analysis and Theory*. Wiley Blackwell.

Szkudlarek E, Brannon EM. (2021). First and Second Graders Successfully Reason About Ratios With Both Dot Arrays and Arabic Numerals. *Child Development*. 92(3):1011-1027. https://doi.org/10.1111/cdev.13470.

Szkudlarek, E., Zhang, H., DeWind, N. K., & Brannon, E. M. (2022). Young children intuitively divide before they recognize the division symbol. *Frontiers in Human Neuroscience*, 16, Article 752190. https://doi.org/10.3389/fnhum.2022.752190

Treisman, A. (1982). Perceptual grouping and attention in visual search for features and for objects. *Journal of Experimental Psychology: Human Perception and Performance*, 8(2), 194–214. https://doi.org/10.1037/0096-1523.8.2.194

Treisman, A. 1986. Features and Objects in Visual Processing. *Scientific American*. 255(5): 114b-25.

von Neumann, J. 1958/2013. *The Computer and the Brain*. Newhaven: Yale University Press.

Wolfe, JM. & Horrowitz, TS. 2017. Five Factors that Guide Attention in Visual Search. *Nature Human Behavior*. 1: 0058

Xu, Y., & Franconeri, S. L. (2015). Capacity for visual features in mental rotation. *Psychological Science*, 26(8), 1241–1251. https://doi.org/10.1177/0956797615585002

Xu, F., & Spelke, E. S. (2000). Large number discrimination in 6-month-old infants. *Cognition*, 74(1), B1–B11. https://doi.org/10.1016/S0010-0277(99)00066-9

Yousif, S. R., & Keil, F. C. (2020). Area, not number, dominates estimates of visual quantities. *Scientific Reports*, 10(1), 13407. https://doi.org/10.1038/s41598-020-68593-z

Yousif, S., & Keil, F. (2021). How we see area and why it matters. *Trends in Cognitive Sciences*, 27(7), 554–557. https://doi.org/10.1016/j.tics.2021.03.017

Zosh, J. M., Halberda, J., & Feigenson, L. (2011). Memory for multiple visual ensembles in infancy. *Journal of Experimental Psychology: General*, 140(2), 141–158. https://doi.org/10.1037/a0022925

How to cite this article: Clarke, S. (2023). "Compositionality and constituent structure in the analogue mind." *Philosophical Perspectives*, 1-29. https://doi.org/10.1111/phpe.12182